

Nonparametric Estimation of Trend in Directional Data

Rudolf Beran¹

*Department of Statistics, University of California, Davis,
One Shields Avenue, Davis CA 95616-8705, USA*

Abstract

Consider measured positions of the paleomagnetic north pole over time. Each measured position may be viewed as a direction, expressed as a unit vector in three dimensions and incorporating some error. In this sequence, the true directions are expected to be close to one another at nearby times. A simple trend estimator that respects the geometry of the sphere is to compute a running average over the time-ordered observed direction vectors, then normalize these average vectors to unit length. This paper treats a considerably richer class of competing directional trend estimators that respect spherical geometry. The analysis relies on a nonparametric error model for directional data in R^q that imposes no symmetry or other shape restrictions on the error distributions. Good trend estimators are selected by comparing estimated risks of competing estimators under the error model. Uniform laws of large numbers, from empirical process theory, establish when these estimated risks are trustworthy surrogates for the corresponding unknown risks.

Keywords: random directions, directional trend model, projected linear estimator, uniform law of large numbers, minimizing estimated risk

1. Introduction

1.1. Preliminaries

Consider measurements on the position of the Earth's north magnetic pole, derived from rock samples collected at various sites. Each observed

Email address: rjberan@ucdavis.edu (Rudolf Beran)

¹This research was supported in part by National Science Foundation Grant DMS-1127914 to the Statistics and Applied Mathematical Sciences Institute, through the SAMSI program on Low-Dimensional Structure in High-Dimensional Systems.

position, usually reported as latitude and longitude, may be represented as a unit vector in R^3 that specifies the direction from the center of the Earth to the point on the Earth's surface with that latitude and longitude. Associated with each such direction vector is the geological dating of the corresponding rock sample. Substantial measurement errors are to be expected in the data. The problem is to extract trend in the position of the north magnetic pole as function of time.

Consider in R^3 the orthonormal basis (j_1, j_2, j_3) in which j_3 is the unit vector pointing to the Earth's geographical north pole and j_1 is the unit vector orthogonal to j_3 that points to longitude 0. Relative to this basis, an observed direction has polar coordinates (θ, ϕ) . Here $\theta \in [0, \pi]$ is the angle, in radians, between j_3 and the observed direction. The angle $\phi \in [0, 2\pi)$ specifies, in radians, the counterclockwise rotation angle in the j_1 - j_2 plane from j_1 to the longitude of the observed direction.

Relative to the same basis, the unit vector with polar coordinates (θ, ϕ) has Cartesian coordinates

$$x_1 = \sin(\theta) \cos(\phi), \quad x_2 = \sin(\theta) \sin(\phi), \quad x_3 = \cos(\theta). \quad (1.1)$$

Cartesian coordinates prove useful in defining trend estimators that operate on directional data. From the Cartesian coordinates of a direction, whether observed or fitted, the polar coordinates may be recovered as

$$(\theta, \phi) = \begin{cases} (\arccos(x_3), \text{atan2}(x_2, x_1)) & \text{if } \text{atan2}(x_2, x_1) \geq 0 \\ (\arccos(x_3), \text{atan2}(x_2, x_1) + 2\pi) & \text{otherwise.} \end{cases} \quad (1.2)$$

The function atan2 , which has domain $R^2 - \{0, 0\}$ and range $(-\pi, \pi]$ is defined by

$$\text{atan2}(v, u) = \begin{cases} \arctan(v/u) & \text{if } u > 0 \\ \arctan(v/u) + \pi & \text{if } u < 0, v \geq 0 \\ \arctan(v/u) - \pi & \text{if } u < 0, v < 0 \\ \pi/2 & \text{if } u = 0, v > 0 \\ -\pi/2 & \text{if } u = 0, v < 0. \end{cases} \quad (1.3)$$

Mathematical programming languages generally provide this function.

Important for visualizing directions in R^3 is the Lambert azimuthal projection of a hemisphere. If a direction has polar coordinates (θ, ϕ) , let

$$(\rho, \psi) = \begin{cases} (2 \sin(\theta/2), \phi) & \text{if } \theta \leq \pi/2 \\ (2 \sin((\pi - \theta)/2), \phi) & \text{if } \theta > \pi/2. \end{cases} \quad (1.4)$$

Then plot the direction as the projected point $(\rho \cos(\psi), \rho \sin(\psi))$ in R^2 , using different plotting symbols according to whether $\theta \leq \pi/2$ (the northern hemisphere) or $\theta > \pi/2$ (the southern hemisphere). The first case, where $\theta \leq \pi/2$, gives an area preserving projection of the northern hemisphere into a disk of radius $\sqrt{2}$. The center of the disk represents the north pole and the perimeter of the disk corresponds to the equator. The second case does likewise for the southern hemisphere, the center of the disk now representing the south pole. See Watson (1983) for a brief derivation of the Lambert (or equal area) projection and a discussion of its use in directional statistics.

Jupp and Kent (1987, pp. 42–45) reported 25 positions of the paleomagnetic north pole, measured in rock specimens from various sites in Antarctica. Each rock specimen was dated, so that the time sequence of the measured positions is known. The left-hand Lambert plot in Fig. 1 displays the measured magnetic pole positions, expressed in polar coordinates and plotted according to (1.4). Line segments join positions adjacent in time. Little pattern emerges. Insight is gained by: (a) using (1.1) to express each observed direction as a unit vector in Cartesian coordinates; (b) forming a 3-year running average of these direction vectors; (c) rescaling each average vector to unit length. Reflection handles the end points in the sequence. Example 1 in subsection 1.2 gives details. After transformation (1.2) back to polar coordinates, the right-hand Lambert plot in Fig. 1 joins three-year running average directions adjacent in time with line segments. Evident now is an apparent convergence, as time goes on, of the paleomagnetic north pole to the geographic north pole.




Figure 1: The measured positions of the paleomagnetic north magnetic pole and two fits: interpolated raw data and interpolated 3-year moving average direction. The interpolating lines merely indicate the time sequence.

In their analysis of the data, Jupp and Kent (1987, Figs. 1 and 2) unwrapped the sphere and directional data onto a plane, used a continuous cubic spline fit to the planar data, then wrapped this fit back onto the sphere. Their continuous spline fit resembles the linearly interpolated 3-year moving directional average except for a surprising kink near the left edge of the Lambert plot. They remarked that this kink lacks physical significance, not being supported by data, and is an artifact of their methodology. In contrast, the running 3-year directional average avoids distortions caused by projection of the sphere onto a plane.

Does the running 3-year average direction provide an effective estimate of the trend in the paleomagnetic data? Clearly we cannot know from a single data-analysis. In this paper, we assess the performance of the three-year running directional average and of certain competing estimators on hypothetical data described by workable probability models. Because this approach examines performance of competing estimators over a broad range of hypothetical data sets, it can reveal their strengths and weaknesses. However, the relationship between data and probability models for data is not untroubled. Insightful comments on the situation include:

- Doob (1972): “The present formal correctness of mathematical probability only helps indirectly in analyzing real probabilistic phenomena. It is unnecessary to stress to statisticians that the relation between mathematics and these phenomena is still obscure.”
- Kolmogorov (1983): “There arises a problem of finding the reasons for the applicability of the mathematical theory of probability to the phenomena of the real world.”
- Tukey (1980): “In practice, methodologies have no assumptions and deliver no certainties.”

1.2. Outline of the Paper

The model problem is to estimate an underlying trend from an observed sequence (y_1, y_2, \dots, y_p) of random column vectors that take values in the unit sphere $S_q = \{y \in R^q : |y| = 1\}$. Each of these direction vectors is indexed by an ordinal covariate and is measured with random error. Mean directions at nearby covariate values may be expected to be close to one another. This paper develops a large class of estimators for the mean directions and supporting theory. Key steps in the development are as follows.

- *Rescaled Symmetric Linear Estimators.* Let Y be the $p \times q$ data matrix with successive rows y'_1, y'_2, \dots, y'_q . Let A be a symmetric $p \times p$ matrix.

For any matrix B , let $\Delta(B) = \text{diag}(BB')$, the diagonal matrix obtained by replacing every off-diagonal entry in BB' with zero. Note that the matrix $\Delta^{-1/2}(B)B$ rescales each row of B to unit length. Consider the estimators

$$\hat{M}(A) = AY, \quad \hat{D}(A) = \Delta^{-1/2}(AY)AY. \quad (1.5)$$

Under the probability model for Y to be described later, $\hat{M}(A)$ is an estimator, usually biased, for the mean matrix $M = E(Y)$. The rows of $\hat{D}(A)$, having unit length, are the implied estimators of the directional means.

Example 1. Let

$$A = \begin{pmatrix} 2/3 & 1/3 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1/3 & 1/3 & 1/3 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1/3 & 1/3 & 1/3 \\ 0 & 0 & 0 & 0 & \dots & 0 & 1/3 & 2/3 \end{pmatrix} \quad (1.6)$$

Then $\hat{D}(A)$ is the 3-year running directional average described in the previous subsection. Evidently, the construction generalizes to weighted running averages of any odd-numbered span. Section 3 treats these and further classes of matrices A useful for estimation of directional trend.

- *Selection Procedure for A .* Section 2 presents a general nonparametric probability model for the matrix Y of directional measurements. The performance of directional mean estimator $\hat{D}(A)$ will be assessed through its risk computed under this model. Let $|\cdot|$ denote the Frobenius matrix norm: $|B|^2 = \text{tr}(BB') = \text{tr}(B'B)$. The *extrinsic loss* of $\hat{D}(A)$ is defined to be

$$L(\hat{D}(A), M) = p^{-1}|\hat{M}(A) - M|^2. \quad (1.7)$$

This choice of loss is algebraically tractable for subsequent theory and will be seen to work reasonably on artificial and real directional data. Because the corresponding *extrinsic risk* of $\hat{D}(A)$ depends on unknown parameters, it is necessary to estimate it from Y . As will be seen in section 2, the estimated risk can be expressed as

$$\hat{R}(A) = p^{-1}[|Y - AY|^2 + (2 \text{tr}(A) - p)\hat{\gamma}^2], \quad (1.8)$$

where $\hat{\gamma}^2$ is an estimator of a dispersion parameter γ^2 to be defined in (2.5). Equation (1.8), for which an asymptotic justification will be provided, is related to some considerations in Mallows (1973).

Candidate estimators $\{\hat{D}(A): A \in \mathcal{A}\}$ are compared through their estimated risks. Of particular potential is the candidate estimator with smallest estimated risk:

$$\hat{D}(\hat{A}), \quad \text{where } \hat{A} = \underset{A \in \mathcal{A}}{\operatorname{argmin}} \hat{R}(A). \quad (1.9)$$

- *Role of Uniform Laws of Large Numbers.* Suppose that \mathcal{A} is a rich class of $p \times p$ symmetric matrices constructed to respect prior vague hypotheses about the mean directions. Suppose that \tilde{A} minimizes the risk of $\hat{D}(A)$ over all $A \in \mathcal{A}$. When does the risk of the adaptive directional estimator $\hat{D}(\hat{A})$ defined above approximate closely the risk of the ideal unrealizable estimator $\hat{D}(\tilde{A})$?

An asymptotic answer, in which the number p of directions observed tends to infinity, draws on uniform laws of large numbers. Let $W(A)$ denote either the loss or the estimated risk $\hat{R}(A)$ of directional estimator $\hat{D}(A)$. Suppose that, as p increases, $W(A)$ converges in L_1 norm to the risk of $\hat{D}(A)$, *uniformly over all* $A \in \mathcal{A}$. Then the risk of the adaptive estimator $\hat{D}(\hat{A})$ converges to the risk of ideal estimator $\hat{D}(\tilde{A})$. Moreover, the plug-in risk estimator $\hat{R}(\hat{A})$ also converges to the risk of $\hat{D}(\hat{A})$. These convergences support use of estimated risk to score the adequacy of the candidate estimator class considered. Section 3 presents theorems of this type, in which empirical process theory plays a central role. Section 5 gives all theorem proofs.

- *Numerical Experiments.* Examples in Section 4 illustrate how vague prior hypotheses about mean directions combined with estimator selection through estimated risks can generate effective estimators of mean directions. The experiments treat three artificial data sets where the truth is known, then analyze further the paleomagnetic north pole data.

1.3. Perspectives

Other directional trend estimators proposed by Watson (1985), Fisher and Lewis (1987), and Jupp and Kent (1987) are curve fitting procedures that rely on analogs of cubic-spline or kernel methods in Euclidean spaces. Important in such treatments are assumptions on the continuity and smoothness of the unknown trend. The estimators of this paper neither fit nor

assume an underlying smooth trend. They are discrete estimators of discrete mean directions that take advantage of any underlying slow variation to reduce estimation risk and perform accountably in any case.

Running average directions are a special case of a more general concept, running extrinsic Fréchet averages. For definitions and some applications of Fréchet means and averages, see Bhattacharya and Patrangenaru (2014). The developments in this paper are potentially extensible to estimating trends in Fréchet means through running Fréchet averages and, more generally, through analogs of (1.5).

For instance, representing axes as projection matrices of rank one enables construction of running average axes. An average axis is defined to be the rank-one projection matrix onto the eigenvector associated with the largest eigenvalue of the average of the observed projection matrices. This definition is simply the Euclidean projection of the averaged rank-one projection matrices into the space of rank-one projection matrices. See Beran and Fisher (1998) for details. A richer class of candidate estimators for trend in axes is obtained by modifying (1.5) so that each row of Y is now a vectorized observed rank-one projection. Each row of $\hat{M}(A) = AY$, reassembled into a symmetric matrix, is then projected, in Euclidean metric, into the set of projection matrices of rank one. The result defines $\hat{D}(A)$ for axes.

2. Nonparametric Model, Risks, and Estimated Risks

This section introduces a nonparametric model for the matrix Y of observed directions in R^q . The quadratic risk of candidate estimator $\hat{D}(A)$ is calculated under this model. Estimated risk is then derived.

2.1. Data Model

Let $\{e_i: 1 \leq i \leq p\}$ be independent identically distributed random unit vectors in S_q such that

$$E(e_i) = \lambda \mu_0, \text{ where } \mu_0 = (1, 0, 0, \dots, 0)' \text{ and } \lambda \neq 0. \quad (2.1)$$

The distribution of e_i and the value of λ are otherwise unknown. No shape assumptions are made on the former. The covariance matrix of e_i then satisfies

$$\Sigma = \text{Cov}(e_i) = E(e_i e_i') - \lambda^2 \mu_0 \mu_0', \quad \gamma^2 = \text{tr}(\Sigma) = 1 - \lambda^2 \geq 0. \quad (2.2)$$

Thus, $0 < |\lambda| \leq 1$, in agreement with geometrical intuition.

The data model specifies we observe $y_i = O_i e_i$ for $1 \leq i \leq p$, where each O_i is an unknown nonrandom $q \times q$ orthogonal matrix. The $\{y_i\}$ are independent random unit vectors in S_q with mean vectors and covariance matrices that depend on i :

$$m_i = E(y_i) = \lambda O_i \mu_0, \quad \text{Cov}(y_i) = O_i \Sigma O_i'. \quad (2.3)$$

The unit vector $\mu_i = O_i \mu_0$ is the *mean direction* of y_i . Because $\lambda \neq 0$, $\mu_i = m_i / |m_i|$.

The $p \times q$ data matrix $Y = \{y_{ij}\}$ has rows y'_1, y'_2, \dots, y'_p , each of unit length, and columns $y_{(1)}, y_{(2)}, \dots, y_{(q)}$. The corresponding *mean matrix* $M = E(Y)$ has rows m'_1, m'_2, \dots, m'_p and columns $m_{(1)}, m_{(2)}, \dots, m_{(q)}$. In particular, $m_{(j)} = E y_{(j)}$. Because the elements of $y_{(j)}$ are independent random variables,

$$\text{Cov}(y_{(j)}) = \text{diag}\{\text{Var}(y_{ij}) : 1 \leq i \leq p\}. \quad (2.4)$$

In view of (2.3),

$$\sum_{j=1}^q \text{Var}(y_{ij}) = \text{tr}(\text{Cov}(y_i)) = \text{tr}(O_i \Sigma O_i') = \text{tr}(\Sigma) = \gamma^2, \quad 1 \leq i \leq p. \quad (2.5)$$

Equations (2.4) and (2.5) imply

$$\sum_{j=1}^q \text{Cov}(y_{(j)}) = \gamma^2 I_p, \quad E\left(\sum_{j=1}^q y_{(j)} y'_{(j)}\right) = \sum_{j=1}^q m_{(j)} m'_{(j)} + \gamma^2 I_p. \quad (2.6)$$

The calculations of risk and quadratic risk in the next subsection draw on these expressions.

2.2. Loss and Risk

For the candidate estimator $\hat{D}(A)$ defined in (1.5), the quadratic *loss* and *risk* are taken to be

$$L(\hat{D}(A), M) = p^{-1} |AY - M|^2, \quad R(\hat{D}(A), M, \gamma^2) = EL(\hat{D}(A), M). \quad (2.7)$$

The next theorem shows how the risk depends on the model only through M and γ^2 .

Theorem 1. *Under the nonparametric data model of section 2, the risk of the candidate estimator \hat{D} is*

$$R(\hat{D}(A), M, \gamma^2) = p^{-1} [\gamma^2 \text{tr}(A^2) + \text{tr}((I_p - A)^2 M M')]. \quad (2.8)$$

Suppose that the $p \times p$ symmetric matrix A has spectral representation

$$A = \sum_{k=1}^s a_k P_k, \quad s \leq p, \quad (2.9)$$

where the $\{a_k\}$ are the distinct eigenvalues and the $\{P_k\}$ are the corresponding mutually orthogonal eigenprojections. Then

$$\begin{aligned} R(\hat{D}(A), M, \gamma^2) &= \sum_{k=1}^s [a_k^2 \tau_k + (1 - a_k)^2 w_k] \\ &= \sum_{k=1}^s [(a_k - \tilde{a}_k)^2 (\tau_k + w_k) + \tau_k \tilde{a}_k], \end{aligned} \quad (2.10)$$

where $\tau_k = p^{-1} \gamma^2 \text{tr}(P_k)$, $w_k = p^{-1} |P_k M|^2$, and $\tilde{a}_k = w_k / (\tau_k + w_k)$. For fixed $\{P_k : 1 \leq k \leq s\}$, the risk $R(\hat{D}(A), M, \gamma^2)$ is minimized when $a_k = \tilde{a}_k$ for every k .

Because $\tilde{a} \in [0, 1]$, it follows that only symmetric matrices A whose eigenvalues all lie in $[0, 1]$ need to be considered when seeking to minimize the risk of $\hat{D}(A)$. Such matrices accomplish Stein-type oracle shrinkage in each of the eigenspaces.

2.3. Estimated Risk

Estimating the risk $R(\hat{D}(A), M, \gamma^2)$ in (2.8) requires estimating γ^2 and $MM' = \sum_{j=1}^q m_{(j)} m'_{(j)}$. It follows from (2.6) that $E(Y Y') = MM' + \gamma^2 I_p$. Consequently the naive estimator $p^{-1} \text{tr}((I_p - A)^2 Y Y')$ for the risk term $p^{-1} \text{tr}((I_p - A)^2 MM')$ in (2.8) is biased upward by the amount $p^{-1} \text{tr}((I_p - A)^2 \gamma^2)$. In general, the bias does not vanish as p increases.

For γ^2 , consider the estimator

$$\hat{\gamma}^2 = [2(p-1)]^{-1} \sum_{i=2}^2 |y_i - y_{i-1}|^2. \quad (2.11)$$

It will be seen in section 3 that, as p increases, $\hat{\gamma}^2$ converges in L_1 norm to γ^2 provided $p^{-1} \sum_{i=2}^2 |m_i - m_{i-1}|^2$ tends to zero. This makes $\hat{\gamma}^2$ a reasonable estimator of γ^2 when successive mean directions vary slowly. Other estimators of γ^2 may be constructed by using higher order differences of the $\{y_i\}$ or taking advantage of replication in the data set. Given a consistent estimator $\hat{\gamma}^2$ of γ^2 , the obvious bias correction yields the *estimated risk*

$$\hat{R}(A) = p^{-1} [\hat{\gamma}^2 \text{tr}(A^2) + \text{tr}((I_p - A)^2 (Y Y' - \hat{\gamma}^2 I_p))] \quad (2.12)$$

Theorem 2. *The estimated risk (2.12) of $\hat{D}(A)$ has the alternative expression*

$$\hat{R}(A) = p^{-1}[|Y - AY|^2 + (2 \operatorname{tr}(A) - p)\hat{\gamma}^2]. \quad (2.13)$$

Suppose the symmetric matrix A has the spectral representation (2.9). Then

$$\begin{aligned} \hat{R}(A) &= \sum_{k=1}^s [a_k^2 \hat{\tau}_k + (1 - a_k)^2 \hat{w}_k] \\ &= \sum_{\hat{w}_k \geq 0} [(a_k - \hat{a}_k)^2 (\hat{\tau}_k + \hat{w}_k) + \hat{\tau}_k \hat{a}_k] + \sum_{\hat{w}_k < 0} [a_k^2 \hat{\tau}_k + (1 - a_k)^2 \hat{w}_k], \end{aligned} \quad (2.14)$$

where $\hat{\tau}_k = p^{-1} \hat{\gamma}^2 \operatorname{tr}(P_k)$, $\hat{w}_k = p^{-1} |P_k Y|^2 - \hat{\tau}_k^2$, and $\hat{a}_k = \hat{w}_k / (\hat{\tau}_k + \hat{w}_k)$ if $\hat{w}_k \geq 0$ but $\hat{a}_k = 0$ if $\hat{w}_k < 0$. For fixed $\{P_k: 1 \leq k \leq s\}$, the estimated risk $\hat{R}(A)$ is minimized when $a_k = \hat{a}_k$ for every k .

The differences between the definitions of \hat{a}_k here and of \tilde{a}_k in the preceding theorem arise because, unlike w_k , the estimated quantity \hat{w}_k can take on negative values.

For a fixed set of $\{P_k: 1 \leq k \leq s\}$, the symmetric matrix $\hat{A} = \sum_{k=1}^s \hat{a}_k P_k$ minimizes estimated risk over the candidate class of matrices A defined by (2.9). The oracle symmetric matrix $\tilde{A} = \sum_{k=1}^s \tilde{a}_k P_k$ minimizes risk over the same candidate class of A . The trustworthiness of \hat{A} as a surrogate for the unrealizable \tilde{A} requires investigation. Key is that the estimated risk $\hat{R}(A)$ converge to the risk $R(\hat{D}(A), M, \gamma^2)$ uniformly over the class \mathcal{A} of symmetric matrices A under consideration as p tends to infinity. This is a question in empirical process theory, addressed in Section 3. For an example of a simpler situation where uniform convergence fails and \hat{A} specifies a poor estimator, see Remark A on p. 1829 of Beran and Dümbgen (1988).

The estimated risk $\hat{R}(A)$ can be negative. This does not affect the rationale for ranking competing estimators according to the order of their estimated risks. Adding a very large constant to the loss function reduces the chances that estimated risk is negative without changing anything essential.

3. Uniform Convergence of Estimated Risks and Adaptation

Uniform laws of large numbers for estimated risks are the main theme of this section. These yield an asymptotic justification for adaptive estimation of mean direction by minimizing estimated risk. Consistent estimation of the dispersion parameter γ^2 is treated.

Many quantities arising in the discussion depend on p . For readability, this dependence is usually suppressed in the notation.

3.1. Assumptions

For any matrix B , let $|B|_{sp} = \sup_{x \neq 0} [|Bx|/|x|]$ denote its spectral norm. The candidate estimators of mean direction are $\{\hat{D}(A): A \in \mathcal{A}\}$, as defined in (1.5). The following assumptions are made:

- $\mathcal{A} = \{A(t): t \in [0, 1]^k, A(t) \text{ is } p \times p \text{ symmetric}\}$.
- $A(t)$ is continuous on $[0, 1]^k$ with $\sup_p \sup_{t \in [0, 1]^k} |A(t)|_{sp} < \infty$.
- $A(t)$ is differentiable on the interior of $[0, 1]^k$, with partial derivatives $\{\nabla_i A(t) = \partial A(t)/\partial t_i: 1 \leq i \leq k\}$. The partial derivatives satisfy $\sup_{p,i} \sup_{t \in [0, 1]^k} |\nabla_i A(t)|_{sp} < \infty$.
- The data matrix Y satisfies the probability model described in subsection 2.1.
- The estimator $\hat{\gamma}^2$ of γ^2 is uniformly L_1 consistent:

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |\hat{\gamma}^2 - \gamma^2| = 0. \quad (3.1)$$

Candidate estimators that meet the assumptions on \mathcal{A} arise naturally. Two examples illustrate:

Example 1 continued. *Running weighted average directions* The span-3 running average directions of (1.6) can be generalized to span-3 running weighted average directions, defined by the class of candidate estimators generated by

$$A(t) = \begin{pmatrix} t_1 + t_2 & t_2 & 0 & 0 & \dots & 0 & 0 & 0 \\ t_2 & t_1 & t_2 & 0 & \dots & 0 & 0 & 0 \\ 0 & t_2 & t_1 & t_2 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & t_2 & t_1 & t_2 \\ 0 & 0 & 0 & 0 & \dots & 0 & t_2 & t_1 + t_2 \end{pmatrix} \quad (3.2)$$

for $t = (t_1, t_2) \in [0, 1]^2$ such that $t_1 + 2t_2 = 1$. The first and last rows are obtained by reflection at the boundary. Because $|A(t)x|^2 \leq [(t_1 + t_2)^2 + t_2^2]|x|^2$, it follows that $|A(t)|_{sp} \leq 5^{1/2}$ for every $t \in [0, 1]^2$. Symmetrically weighted running average directions were already used by Irving (1977) in

a data analysis of continental drift data. Note that $\nabla_1 A(t) = I_p$ and

$$\nabla_2 A(t) = \begin{pmatrix} 1 & 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 & 1 \end{pmatrix}. \quad (3.3)$$

Thus, $\sup_{t \in [0,1]^2} |\nabla_1 A(t)|_{sp} = 1$ and $\sup_{t \in [0,1]^2} |\nabla_2 A(t)|_{sp} \leq 2^{1/2}$.

Example 2. *Multiply penalized least squares.* Penalized least squares estimators for Euclidean means inspire the following development for directional means. Let $t = \{t_i: 1 \leq i \leq k\}$ be penalty weights such that $0 \leq t_i \leq 1$. Let c be a positive constant, possibly very large. Let $\{Q_i: 1 \leq i \leq k\}$ be symmetric, positive semi-definite matrices, normalized so that $|Q_i|_{sp} = 1$ for $1 \leq i \leq k$. Consider the class of candidate directional mean $\hat{D}(A(t))$ estimators generated by

$$A(t) = (I_p + Q(t))^{-1}, \quad \text{where } Q(t) = c \sum_{i=1}^k t_i Q_i, \quad t \in [0, 1]^k. \quad (3.4)$$

Note that $A(t)Y$ is the value of M that minimizes the penalized least squares (PLS) criterion $|Y - M|^2 + \text{tr}(M'Q(t)M)$.

Let $B(t) = A^{-1}(t)$. The eigenvalues of $B(t)$ all lie in $[1, 1 + ck]$ for every $t \in [0, 1]^k$. Consequently, $\sup_{t \in [0,1]^k} |A(t)|_{sp} \leq 1$. Moreover

$$\nabla_i A(t) = -B^{-1}(t)[\nabla_i B(t)]B^{-1}(t) = cB^{-1}(t)Q_i B^{-1}(t). \quad (3.5)$$

Thus,

$$|\nabla_i A(t)|_{sp} \leq c|B^{-1}(t)|_{sp}|Q_i|_{sp}|B^{-1}(t)|_{sp} \leq c(1 + ck)^{-2} \quad \forall t \in [0, 1]^k. \quad (3.6)$$

3.2. Uniform Laws of Large Numbers and Adaptation

The uniform laws of large numbers in this section hold for matrices in the class $\mathcal{A} = \{A(t): t \in [0, 1]^k\}$ satisfying the assumptions of the preceding subsection.

Theorem 3. *Suppose that the assumptions of subsection 3.1 hold. Let $W(A)$ denote either the loss $L(\hat{D}(A), M)$ or the estimated risk $\hat{R}(A)$ of candidate directional estimator $\hat{D}(A)$. Then, for every finite $\gamma^2 > 0$,*

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E}[\sup_{A \in \mathcal{A}} |W(A) - R(\hat{D}(A), M, \gamma^2)|] = 0. \quad (3.7)$$

Theorem 3 asserts that the loss, risk, and estimated risk of candidate estimator $\hat{D}(A)$ converge together, uniformly over all $A \in \mathcal{A}$, as the number p of mean directions observed increases. As in the Introduction, let

$$\hat{A} = \operatorname{argmin}_{A \in \mathcal{A}} \hat{R}(A), \quad \tilde{A} = \operatorname{argmin}_{A \in \mathcal{A}} R(\hat{D}(A), M, \gamma^2). \quad (3.8)$$

The uniform convergence 3.7 justifies using the *adaptive* directional estimator $\hat{D}(\hat{A})$ as a surrogate for the unrealizable oracle estimator $\hat{D}(\tilde{A})$, which achieves minimal risk over the candidate class $\{\hat{D}(A) : A \in \mathcal{A}\}$.

Theorem 4. *Suppose that the assumptions of subsection 3.1 hold. Then, for every finite $\gamma^2 > 0$,*

$$\lim_{p \rightarrow \infty} \sup_M |R(\hat{D}(\hat{A}), M, \gamma^2) - R(\hat{D}(\tilde{A}), M, \gamma^2)| = 0. \quad (3.9)$$

Moreover, for V equal to either the loss or risk of the adaptive directional estimator $\hat{D}(\hat{A})$,

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |\hat{R}(\hat{A}) - V| = 0. \quad (3.10)$$

Theorem 4 shows that the risk of the adaptive directional estimator $\hat{D}(\hat{A})$ converges to the best possible risk achievable over the class of candidate estimators $\{\hat{D}(A) : A \in \mathcal{A}\}$. Equation (3.10) establishes that the plug-in estimated risk $\hat{R}(\hat{A})$ converges to the loss or risk of $\hat{D}(\hat{A})$. This makes estimated risk a credible metric for scoring the relative performance of competing candidate directional estimators.

3.3. Consistent Estimation of γ^2

In the situation treated in this paper, where there one noisy observation on each mean direction, estimating γ^2 consistently requires further model additional assumptions. The next theorem shows how the presence of slow variation between successive unknown mean directions can be used to estimate γ^2 . Analogous results hold for estimators of γ^2 based on higher order differences of the observed directions.

Theorem 5. *Suppose that the assumptions of subsection 3.1 prior to (3.1) hold and that*

$$\lim_{p \rightarrow \infty} [2(p-1)]^{-1} \sum_{i=2}^p |m_i - m_{i-1}|^2 = 0. \quad (3.11)$$

Let

$$\hat{\gamma}^2 = [2(p-1)]^{-1} \sum_{i=2}^p |y_i - y_{i-1}|^2. \quad (3.12)$$

Then, for every finite $\gamma^2 > 0$,

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |\hat{\gamma}^2 - \gamma^2| = 0. \quad (3.13)$$

4. Numerical Experiments

This section compares, on artificial directional data in three dimensions, two natural adaptive penalized least squares (PLS) directional mean estimators and the span-3 running directional average described in the Introduction. The adaptive PLS estimators are then applied to the paleomagnetic north pole data described in the Introduction.

4.1. Candidate k -th difference PLS Estimators

Consider the $(g-1) \times g$ first-difference matrix $\text{Dif}(g) = \{\delta_{u,v}\}$ in which $\delta_{u,u} = 1$, $\delta_{u,u+1} = -1$ for every u and all other elements are 0. The $(p-1) \times p$ first-difference matrix is $\Delta_1 = \text{Dif}(p)$. Define the $(p-d) \times p$ d -th difference matrix recursively through

$$\Delta_1 = \text{Dif}(p), \quad \Delta_d = \text{Dif}(p-d+1)\Delta_{d-1}, \quad 2 \leq d \leq p-1 \quad (4.1)$$

The d -th difference PLS directional mean estimator $\hat{D}(A(t))$ is defined by a special case of equation (3.4), in which

$$A(t) = (I_p + ct\Delta'_d\Delta_d/|\Delta'_d\Delta_d|_{sp})^{-1}, \quad t \in [0, 1]. \quad (4.2)$$

This PLS fit nudges the raw data Y towards a smoother fit, according to the choice of the difference-order d in the penalty term and the magnitude of the penalty weight $ct > 0$.

4.2. Generating Artificial Directional Trend Data

Let U_1, U_2 be independent random variables, each uniformly distributed on $[0, 1]$. Let

$$\delta = \log[1 + (\exp(2\kappa) - 1)U_1], \quad \theta = \cos^{-1}(\delta/\kappa - 1), \quad \phi = 2\pi U_2. \quad (4.3)$$

The random unit vector z in R^3 with polar coordinates (θ, ϕ) has a Fisher-Langevin distribution with mean direction $\nu_0 = (0, 0, 1)'$ and precision κ (cf. Mardia and Jupp (2000)). For any unit vector μ , the 3×3 orthogonal matrix

$$\Omega(\mu) = (1 + \nu'_0\mu)^{-1}(\nu_0 + \mu)(\nu_0 + \mu)' - I_3 \quad (4.4)$$

rotates ν_0 into μ (cf. Watson (1983, p. 28)). Hence the random vector $\Omega(\mu)z$ has a Fisher-Langevin distribution with mean direction μ and precision κ .

Let $f: [0, 1] \rightarrow [0, \pi]$ and $g: [0, 1] \rightarrow [0, 2\pi]$ be designated functions. The pairs

$$\theta_i = f(i/(p+1)), \quad \phi_i = g(i/(p+1)), \quad 1 \leq i \leq p \quad (4.5)$$

define, in polar coordinates, a sequence of unit vectors $\{\mu_i: 1 \leq i \leq p\}$. Let $\{z_i: 1 \leq i \leq p\}$ be independent random unit vectors, each constructed as in the preceding paragraph to have a Fisher-Langevin distribution with mean direction ν_0 and precision κ . Then the random unit vectors

$$y_i = \Omega(\mu_i)z_i \quad (4.6)$$

are independent and y_i has a Fisher-Langevin distribution with mean direction μ_i and precision κ . This construction is a special case of the nonparametric data model defined in subsection 2.1. The Fisher-Langevin model has a historical record of fitting some small directional data sets plausibly (cf. Fisher, Lewis and Embleton (1985), Mardia and Jupp (2000)).

4.3. Numerical Trials on Artificial Data

Figures 2, 3, and 4 exhibit competing directional trend fits to three artificial data sets constructed as described in subsection 4.2. These are compared with the true mean directions in each case. In each data set, $p = 150$ directional means $\{\mu_i\}$ are observed with error. The observed directions $\{y_i: 1 \leq i \leq 150\}$ are pseudorandom unit vectors such that y_i has a Fisher-Langevin distribution with mean direction μ_i and precision $\kappa = 200$. The functions f and g that define, through (4.5), the polar coordinates for each of the three sets of true directional means $\{\mu_i\}$ are:

- *Wobble*: $f(t) = .3\pi(t + .2 + .15\sin(36\pi t))$ while $g(t) = 4\pi t$.
- *Bat*: $f(t) = .8\pi(t - .5)$ while $g(t) = 4\pi\sin(6\pi t)$.
- *Jumps*: $f(t) = .2\pi$ if $0 \leq t \leq .15$, $= .1\pi$ if $.15 < t \leq .3$, $= .4\pi$ if $.3 < t \leq .45$, $= .2\pi$ if $.45 < t \leq .65$, $= .3\pi$ if $.65 < t \leq .8$, and $= .4\pi$ if $.8 < t \leq 1$ while $g(t) = 2\pi t$.

The artificial Wobble data are inspired by observations on the Chandler-wobble of the geographic north pole, scaled up to wander over a larger part of the northern hemisphere and given greater measurement errors. Brillinger (1973) analysed actual Chandler-wobble data by using time-series methods in the tangent plane of the sphere at the north pole.

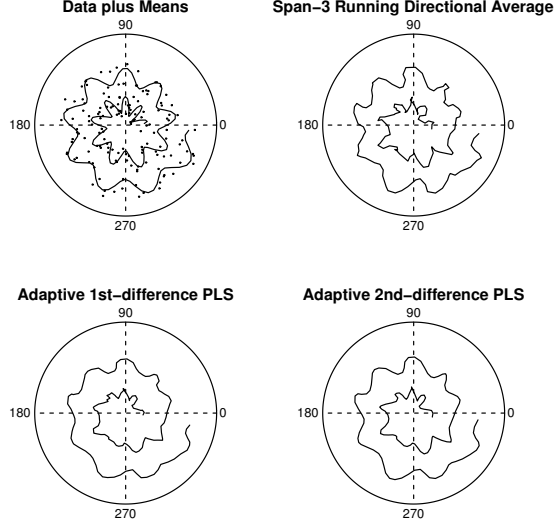


Figure 2: *The Wobbly Spiral artificial data, together with the underlying true mean directions, plus three competing fits to the data. Interpolating lines indicate the time sequence of the true and the estimated mean directions.*

Three fits were computed for each artificial data set: the span-3 running average, the adaptive 1st-difference PLS estimator, and the adaptive 2nd-difference PLS estimator (cf. subsection 4.1). The estimator $\hat{\gamma}^2$ defined in (3.12) was used to compute estimated risks.

Fig. 2 displays the true means and the competing fits to the Wobble data, with linear interpolation added solely to guide the eye. The task is estimation of *discrete* successive mean directions observed with error. From smallest to largest, the estimated risks of the competing directional trend estimators are: .0020 (running average); .0022 (adaptive 2nd-difference PLS); .0032 (adaptive 1st-difference PLS). These relatively similar estimated risks contrast with the much larger estimated risk .0138 for the naive estimator consisting of the observed directions $\{y_i\}$. To the eye, the smoothed fits vary in details but not in major features.

Fig. 3 presents the true means and the competing fits to the Bat data, again with linear interpolation added solely to guide the eye. From smallest to largest, the estimated risks of the competing directional trend estimators are: .0001 (adaptive 2nd-difference PLS); .0018 (adaptive 1st-difference PLS); .0023 (running average). These values contrast with the much larger

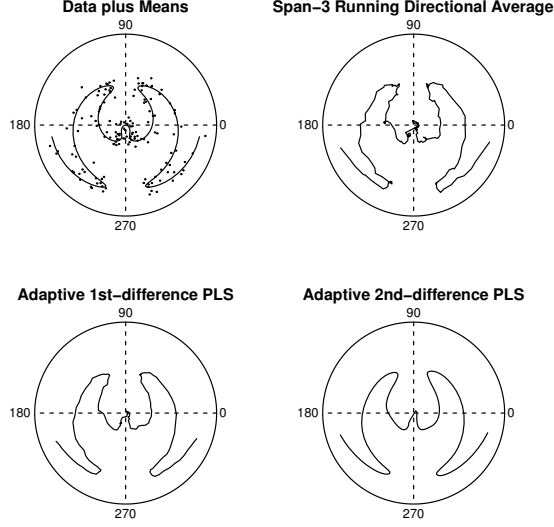


Figure 3: *The Bat artificial data, together with the underlying true mean directions, plus three competing fits to the data. Interpolating lines indicate the time sequence of the true and estimated mean directions.*

estimated risk .0117 for the naive estimator. Visually, the adaptive 2nd-difference PLS estimator is clearly the best while the 1st-difference PLS estimator seems slightly better than the running average.

Fig. 4 gives the true means and the competing fits to the Jumps data. From smallest to largest, the estimated risks of the competing directional trend estimators are: .0028 (adaptive 2nd-difference PLS); .0031 (adaptive 1st-difference PLS); .0038 (running average). These values contrast with the much larger estimated risk .0155 for the naive estimator. To the eye, the 2nd-difference PLS fit exhibits the fewest extraneous artifacts. All three fitting techniques face difficulty in identifying the sharp transitions.

4.4. *Fitting the Paleomagnetic North Pole Data*

Fig. 5 exhibits the adaptive first- and second-difference PLS fits to the sequence of paleomagnetic north pole observations described in the Introduction. Again, linear interpolation has been added merely to guide the eye between successive estimated mean directions. These fits are to be compared with the raw data and with the 3-year running average directions displayed in Fig.1. From smallest to largest, the estimated risks of the competing

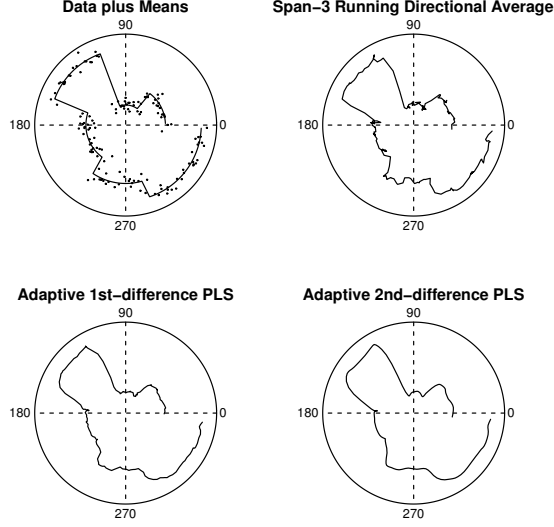


Figure 4: *The Jumps artificial data, together with the underlying true mean directions, plus three competing fits to the data. Interpolating lines indicate the time sequence of the true and the estimated mean directions.*

directional trend estimators are: .0158 (adaptive 2nd-difference PLS); .0230 (adaptive 1st-difference PLS); .0303 (running average). These values contrast with the much larger estimated risk .1055 for the naive estimator consisting of the raw data. The adaptive 2nd-difference PLS estimator—the competitor with smallest estimated risk—is arguably the most pleasing to the eye. It tells a clear story of the paleomagnetic north pole spiraling, with time, to the vicinity of the Earth’s geographic north pole.

The remark of Tukey (1980) on methodologies, quoted in the Introduction, remains pertinent. The fitting procedures treated in this paper are scored credibly by their estimated risks under the nonparametric data model of subsection 2.1—the asymptotic theory and numerical trials in sections 3 and 4 establish this point. The directional trend fits obtained for the paleomagnetic data are highly suggestive scientifically. Nevertheless, no certainties follow. There is no intrinsic reason for observational errors in the paleomagnetic data to obey a probability model.




Figure 5: The paleomagnetic north pole data with the first- and second-difference adaptive PLS estimated mean directions. Interpolating lines indicate the time sequence of the fitted mean directions.

5. Proofs

5.1. Ancillary Results

The theorem proofs draw on standard properties of the spectral and Frobenius matrix norms and on two preliminary Propositions. The first of these, from the theory of weak convergence in $C[0, 1]^k$, is restated here for convenient reference in subsequent argument. Lütkepohl (1997) summarizes pertinent results on the relations among matrix norms.

Proposition 1. *Let the $\{U_p: p \geq 1\}$ be random elements of $C([0, 1]^k)$. Let $\text{plim}_{p \rightarrow \infty}$ denote the limit on probability as $p \rightarrow \infty$. Suppose that*

$$\begin{aligned} \text{plim}_{p \rightarrow \infty} U_p(t) &= 0 \quad \forall t \in [0, 1]^k, \\ \lim_{\delta \rightarrow 0} \limsup_{p \rightarrow \infty} \mathbb{P} \left[\sup_{|t-s| \leq \delta} |U_p(t) - U_p(s)| \geq \epsilon \right] &= 0 \quad \forall \epsilon > 0. \end{aligned} \quad (5.1)$$

Then $\text{plim}_{p \rightarrow \infty} \sup_{t \in [0, 1]^k} |U_p(t)| = 0$.

The next Proposition presents the core result underlying the proof of Theorem 3. The following notation and assumptions are required:

- For $j = 1, 2$, the $p \times p$ symmetric matrices $\{C_j(t): t \in [0, 1]^k\}$ are continuous on $[0, 1]^k$ with $\sup_p \sup_{t \in [0, 1]^k} |C_j(t)|_{sp} < \infty$.

- For $j = 1, 2$, the matrix $C_j(t)$ is differentiable on the interior of $[0, 1]^k$, with partial derivatives $\{\nabla_i C_j(t) = \partial C_j(t)/\partial t_i : 1 \leq i \leq k\}$ such that $\sup_{p,i} \sup_{t \in [0,1]^k} |\nabla_i C_j(t)|_{sp} < \infty$.
- The random vector $z = (z_1, z_2, \dots, z_p)'$ has $E(z) = 0$ and $\text{Cov}(z) = K = \{\text{diag}(k_{ii})\}$. The components of z are independent random variables such that $\sup_{1 \leq i \leq p} |z_i| \leq 1$ w.p.1.

Proposition 2. *Suppose these assumptions hold. Let $b = (b_1, b_2, \dots, b_p)'$ be any constant vector such that $|b_i| \leq 1$ for $1 \leq i \leq p$. Let*

$$U_p(t) = p^{-1}[b'C_1(t)z + \{z'C_2(t)z - \text{tr}(C_2(t)K)\}]. \quad (5.2)$$

Then

$$\lim_{p \rightarrow \infty} E[\sup_{t \in [0,1]^k} |U_p(t)|] = 0. \quad (5.3)$$

Proof. *Pointwise consistency.* For notational simplicity, write C_j for $C_j(t)$ in the calculations. From the assumptions, $E(b'C_1z) = E[z'C_2z - \text{tr}(C_2)K] = 0$, $k_{ii} = \text{Var}(z_i) \leq 1$, $\text{Var}(z_i^2) \leq 1$ and $|b|^2 \leq p$. Consequently,

$$\text{Var}(p^{-1}b'C_1z) = p^{-2}b'C_1KC_1b \leq p^{-2}|b|^2|C_1KC_1|_{sp} \leq p^{-1}|C_1|_{sp}^2. \quad (5.4)$$

Moreover, if $C_2 = \{c_{2,ij}\}$, then $z'C_2z = \sum_i c_{2,ii}z_i^2 + 2\sum_{i < j} c_{2,ij}z_iz_j$ and

$$\text{Var}(p^{-1}z'C_2z) \leq p^{-2}(\sum_i c_{2,ii}^2 + 4\sum_{i < j} c_{2,ij}^2) \leq 2p^{-2}|C_2|^2 \leq 2p^{-1}|C_2|_{sp}^2. \quad (5.5)$$

The upper bounds in (5.4) and (5.5) both tend as zero as p increases. Hence,

$$\text{plim}_{p \rightarrow \infty} U_p(t) = 0 \quad \forall t \in [0, 1]^k. \quad (5.6)$$

Uniform consistency. For every $s, t \in [0, 1]^k$,

$$U_p(s) - U_p(t) = p^{-1} \sum_{i=1}^k (s_i - t_i) [b'\nabla_i C_1z + \{z'\nabla_i C_2z - \text{tr}(\nabla_i C_2K)\}], \quad (5.7)$$

where $\nabla_i C_j = \nabla_i C_j(\bar{s})$ for some \bar{s} on the line segment that joins s and t . Therefore,

$$\sup_{|s-t| \leq \delta} |U_p(s) - U_p(t)| \leq \delta p^{-1} \sum_{i=1}^k [|b'\nabla_i C_1z| + |z'\nabla_i C_2z| + |\text{tr}(\nabla_i C_2K)|]. \quad (5.8)$$

Because $|K^{1/2}|_{sp} \leq 1$, $E|z|^2 \leq p$, and $|b|^2 \leq p$,

$$\begin{aligned} p^{-1}E|b'\nabla_i C_1 z| &\leq p^{-1}|b|E|z||\nabla_i C_1|_{sp} \leq |\nabla_i C_1|_{sp}, \\ p^{-1}E|z'\nabla_i C_2 z| &\leq p^{-1}E|z|^2|\nabla_i C_2|_{sp} \leq |\nabla_i C_2|_{sp}, \\ p^{-1}|\text{tr}(K^{1/2}\nabla_i C_2 K^{1/2})| &\leq |K^{1/2}\nabla_i C_2 K^{1/2}|_{sp} \leq |\nabla_i C_2|_{sp}. \end{aligned} \quad (5.9)$$

In view of (5.8), display (5.9), and the Proposition assumptions, Markov's inequality establishes existence of a finite constant c , not depending on p , such that

$$P\left[\sup_{|s-t|\leq\delta}|U_p(s)-U_p(t)|\geq\epsilon\right]\leq c\delta. \quad (5.10)$$

Consequently,

$$\lim_{\delta\rightarrow 0}\limsup_{p\rightarrow\infty}P\left[\sup_{|t-s|\leq\delta}|U_p(t)-U_p(s)|\geq\epsilon\right]=0. \quad (5.11)$$

Limits (5.6) and (5.11), together with Proposition 1, establish

$$\text{plim}_{p\rightarrow\infty}\sup_{t\in[0,1]^k}|U_p(t)|=0. \quad (5.12)$$

L_1 uniform consistency. Observe that

$$\sup_{t\in[0,1]^k}|U_p(t)|\leq p^{-1}\sup_{t\in[0,1]^k}[|b'C_1(t)z|+|z'C_2(t)z|+|\text{tr}(C_2(t)K)|] \quad (5.13)$$

Because $|b|^2 \leq p$, $|z|^2 \leq p$ w.p.1, and $|K^{1/2}|_{sp} \leq 1$,

$$\begin{aligned} p^{-1}|b'C_1(t)z| &\leq |C_1(t)|_{sp}, \quad p^{-1}|z'C_2(t)z| \leq |C_2(t)|_{sp}, \\ p^{-1}|\text{tr}(C_2(t)K)| &= p^{-1}|\text{tr}(K^{1/2}C_2(t)K^{1/2})| \leq |C_2(t)|_{sp}. \end{aligned} \quad (5.14)$$

Thus, the sequence $\{\sup_{t\in[0,1]^k}|U_p(t)|: p \geq 1\}$ is bounded and (5.12) may be strengthened to the desired L_1 convergence (5.3). \square

5.2. Theorem Proofs

Proof of Theorem 1. In view of (2.6), (2.7), and the symmetry of A , the risk is

$$\begin{aligned} p^{-1}E|AY - M|^2 &= p^{-1}\sum_{j=1}^q E|Ay_{(j)} - m_{(j)}|^2 \\ &= p^{-1}\sum_{j=1}^q [\text{tr}(A^2)\text{Cov}(y_{(j)}) + \text{tr}((I_p - A)^2 m_{(j)} m'_{(j)})] \\ &= p^{-1}[\gamma^2 \text{tr}(A^2) + \text{tr}((I_p - A)^2 M M')]. \end{aligned} \quad (5.15)$$

This establishes (2.8).

The spectral representation (2.9) for A entails

$$\begin{aligned}\gamma^2 \operatorname{tr}(A^2) &= \sum_{k=1}^s \gamma^2 a_k^2 \operatorname{tr}(P_k) \\ \operatorname{tr}((I_p - A)^2 M M') &= \sum_{k=1}^s (1 - a_k)^2 |P_k M|^2.\end{aligned}\quad (5.16)$$

Applying this to (2.8) yields the first expression in (2.10). Completing the square then gives the second line in (2.10). \square

Proof of Theorem 2. The expression (2.12) for $\hat{R}(A)$ entails

$$\begin{aligned}\hat{R}(A) &= p^{-1} [\operatorname{tr}((I_p - A)^2 Y Y') + \hat{\gamma}^2 (\operatorname{tr}(A^2) - \operatorname{tr}((I_p - A)^2))] \\ &= p^{-1} [|Y - AY|^2 + \operatorname{tr}(2A - I_p) \hat{\gamma}^2],\end{aligned}\quad (5.17)$$

which yields the desired expression (2.13).

Arguing as in (5.16) yields

$$\gamma^2 \operatorname{tr}(A^2) = \sum_{k=1}^s \gamma^2 a_k^2 \operatorname{tr}(P_k), \quad \operatorname{tr}((I_p - A)^2 Y Y') = \sum_{k=1}^s (1 - a_k)^2 |P_k Y|^2 \quad (5.18)$$

and

$$\hat{\gamma}^2 \operatorname{tr}((I_p - A)^2) = \sum_{k=1}^s \hat{\gamma}^2 (1 - a_k)^2 \operatorname{tr}(P_k). \quad (5.19)$$

Substituting (5.18) and (5.19) into the first line in (5.17) yields the first expression in (2.14). Completing the square for the summands with $\hat{w}_k \geq 0$ then gives the second line in (2.14). \square

Proof of Theorem 3. The theorem assumptions stated in subsection 3.1 imply the following properties for $B(t) = A^2(t)$ and $\bar{B}(t) = (I_p - A(t))^2$:

$$\begin{aligned}\sup_p \sup_{t \in [0,1]^k} |B(t)|_{sp} &< \infty, & \sup_{p,i} \sup_{t \in [0,1]^k} |\nabla_i B(t)|_{sp} &< \infty, \\ \sup_p \sup_{t \in [0,1]^k} |\bar{B}(t)|_{sp} &< \infty, & \sup_{p,i} \sup_{t \in [0,1]^k} |\nabla_i \bar{B}(t)|_{sp} &< \infty.\end{aligned}\quad (5.20)$$

The bounds on derivatives use the identity $\nabla_i B(t) = \nabla_i A(t) \cdot A(t) + A(t) \cdot \nabla_i A(t)$.

The case $W(A) = \hat{R}(A)$. Define $\check{R}(A)$ by replacing $\hat{\gamma}^2$ with γ^2 in expression (2.12) for $\hat{R}(A)$. For simplicity, the argument t is omitted in the

following calculations, where B, \bar{B} stand for $B(t), \bar{B}(t)$ respectively and so forth. Because

$$\begin{aligned} |\hat{R}(A) - \check{R}(A)| &\leq p^{-1} |\hat{\gamma}^2 - \gamma^2| (|\text{tr}(B)| + |\text{tr}(\bar{B})|) \\ &\leq |\hat{\gamma}^2 - \gamma^2| (|B|_{sp} + |\bar{B}|_{sp}) \end{aligned} \quad (5.21)$$

and $\hat{\gamma}^2$ is L_1 consistent by (3.1), it suffices to prove the case $W(A) = \check{R}(A)$.

The random vector $u_{(j)} = y_{(j)} - m_{(j)}$ has $E(u_{(j)}) = 0$. From (2.4), (2.5) and (2.6), $V_j = \text{Cov}(u_{(j)})$ is a diagonal matrix such that $\sum_{j=1}^q V_j = \gamma^2 I_p$. Because each row of Y and of M is a unit vector, the components of $u_{(j)} = \{u_{ij}\}$ are independent random variables such that $\sup_{1 \leq i \leq p} |u_{ij}| \leq 2$ w.p.1.

Expression (2.8) for $R(\hat{D}(A), M, \gamma^2)$ and the definition above of $\check{R}(A)$ yield

$$\begin{aligned} \check{R}(A) - R(\hat{D}(A), M, \gamma^2) &= p^{-1} \text{tr}[\bar{B} \{ \sum_{j=1}^q (y_{(j)} y'_{(j)} - m_{(j)} m'_{(j)}) - \gamma^2 I_p \}] \\ &= p^{-1} \text{tr}[\bar{B} \sum_{j=1}^q (u_{(j)} u'_{(j)} + m_{(j)} u'_{(j)} + u_{(j)} m'_{(j)} - V_j)] \\ &= p^{-1} \sum_{j=1}^q [2m'_{(j)} \bar{B} u_{(j)} + \{u'_{(j)} \bar{B} u_{(j)} - \text{tr}(\bar{B} V_j)\}]. \end{aligned} \quad (5.22)$$

Apply Proposition 2 separately to each summand on the right side of (5.22) after setting $z = u_{(j)}/2$, $b = m_{(j)}$, $C_1(t) = C_2(t) = 4\bar{B}(t)$, and $K = V_j/4$. This completes the proof of (3.7) when $W(A) = \hat{R}(A)$.

The case $W(A) = L(\hat{D}(A), M)$. The argument is similar. Let $F = A^2 - A$, a symmetric matrix. Note that the loss defined in (2.7) may be rewritten

$$\begin{aligned} L(\hat{D}(A), M) &= p^{-1} \sum_{j=1}^q |A u_{(j)} + (A - I_p) m_{(j)}|^2 \\ &= p^{-1} \sum_{j=1}^q [2m'_{(j)} F u_{(j)} + u'_{(j)} B u_{(j)} + m'_{(j)} \bar{B} m_{(j)}]. \end{aligned} \quad (5.23)$$

On the other hand, the risk (2.8) may be rewritten

$$R(\hat{D}(A), M, \gamma^2) = p^{-1} \sum_{j=1}^q [\text{tr}(B V_j) + m'_{(j)} \bar{B} m_{(j)}]. \quad (5.24)$$

Consequently, $\Delta = L(\hat{D}(A), M) - R(\hat{D}(A), M, \gamma^2)$ satisfies

$$\Delta = p^{-1} \sum_{j=1}^q [2m'_{(j)} F u_{(j)} + \{u'_{(j)} B u_{(j)} - \text{tr}(B V_j)\}]. \quad (5.25)$$

Apply Proposition 2 separately to each summand on the right side of (5.25) after setting $z = u_{(j)}/2$, $b = m_{(j)}$, $C_1(t) = 4(A(t)^2 - A(t))$, $C_2(t) = 4B(t)$, and $K = V_j/4$. This completes the proof of (3.7) when $W(A) = L(\hat{D}(A), M)$. \square

Proof of Theorem 4. Let $r(A, M, \gamma^2) = R(\hat{D}(A), M, \gamma^2)$. It suffices to show that (3.7) in Theorem 3 implies

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |Z - r(\tilde{A}, M, \gamma^2)| = 0, \quad (5.26)$$

where Z is either the loss of $\hat{D}(\hat{A})$ or the estimated risk $\hat{R}(\hat{A})$. The three limits to be proved in (3.9) and (3.10) are then immediate.

First, (3.7) with $W(A) = \hat{R}(A)$ implies

$$\begin{aligned} \lim_{p \rightarrow \infty} \sup_M \mathbb{E} |\hat{R}(\hat{A}) - r(\hat{A}, M, \gamma^2)| &= 0, \\ \lim_{p \rightarrow \infty} \sup_M \mathbb{E} |\hat{R}(\hat{A}) - r(\tilde{A}, M, \gamma^2)| &= 0. \end{aligned} \quad (5.27)$$

Consequently, (5.26) holds for $Z = \hat{R}(\hat{A})$ and

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |r(\hat{A}, M, \gamma^2) - r(\tilde{A}, M, \gamma^2)| = 0. \quad (5.28)$$

Second, (3.7) with $W(A) = L(\hat{D}(A), M)$ yields

$$\lim_{p \rightarrow \infty} \sup_M \mathbb{E} |L(\hat{D}(\hat{A}), M) - r(\hat{A}, M, \gamma^2)| = 0. \quad (5.29)$$

This and (5.28) establish that (5.26) holds for $Z = L(\hat{D}(\hat{A}), M)$. \square

Proof of Theorem 5. By algebraic expansion of (3.12),

$$\hat{\gamma}^2 = [2(p-1)]^{-1} \sum_{i=2}^p [|y_i - m_i|^2 + |y_{i-1} - m_{i-1}|^2] + |m_i - m_{i-1}|^2 + T_p, \quad (5.30)$$

where $(p-1)T_p$ is the quantity

$$\sum_{i=2}^p [(y_i - m_i)'(m_i - m_{i-1}) - (y_{i-1} - m_{i-1})'(m_i - m_{i-1}) - (y_i - m_i)'(y_{i-1} - m_{i-1})']. \quad (5.31)$$

Recall from Section 2 that $|y_i| = |m_i| = 1$, $E(y_i) = m_i$, and $E|y_i - m_i|^2 = \text{tr}(\Sigma) = \gamma^2$ for each i . Straightforward variance calculations plus Chebyshev's inequality establish that

$$\begin{aligned} \text{plim}_{p \rightarrow \infty} (p-1)^{-1} \sum_{i=2}^p |y_i - m_i|^2 &= \gamma^2 = \text{plim}_{p \rightarrow \infty} (p-1)^{-1} \sum_{i=2}^p |y_{i-1} - m_{i-1}|^2 \\ \text{plim}_{p \rightarrow \infty} T_p &= 0. \end{aligned} \quad (5.32)$$

This together with the assumption (3.11) establishes $\text{plim}_{p \rightarrow \infty} \hat{\gamma}^2 = \gamma^2$. This may be strengthened to L_1 convergence because each of the normalized sums above is bounded. A closer look at the argument verifies the desired uniformity in M . \square

References

- Beran, R., Dümbgen, L., 1998. Modulation of estimators and confidence sets. *Annals of Statistics* 26, 1826–1856.
- Beran, R., Fisher, N.I., 1998. Nonparametric comparison of mean directions or mean axes. *Annals of Statistics* 26, 472–493.
- Bhattacharya, R., Patrangenaru, V., 2014. Statistics on manifolds and landmarks based image analysis: a nonparametric theory with applications. *Journal of Statistical Planning and Inference* 145, 1–22.
- Brillinger, D.R., 1973. An empirical investigation of the Chandler wobble and two proposed excitation processes. *Bulletin of the International Statistical Society*, Book3 413–433.
- Doob, J.L., 1972. William Feller and Twentieth Century Probability. In: *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability* (Le Cam, L.M., Neyman, J., Scott, E.L. Eds.), University of California Press, pp. xv–xx.
- Fisher, N.I., Lewis, T., Embleton, B.J.J., 1985. *Statistical Analysis of Spherical Data*. Cambridge University Press.
- Fisher, N.I., Lewis, T., 1987. A note on spherical splines. *Journal of the Royal Statistical Society, Series B* 47, 482–488.
- Irving, E., 1977. Drift of the major continental blocks since the Devonian. *Nature* 270, 304–309.

- Jupp, P.E., Kent, J., 1987. Fitting smooth paths to spherical data. *Applied Statistics* 36, 34–46.
- Kolmogorov, A.N., 1983. On logical foundations of probability theory. In: *Probability Theory and Mathematical Statistics* (Itô, K., Prokhorov, J.V. Eds.), *Probability Theory and Mathematical Statistics*, vol. 1021, Springer, pp. 1–5.
- Lütkepohl, H., 1997. *Handbook of Matrices*. Wiley.
- Mallows, C., 1973. Some comments on C_p . *Technometrics* 15, 661–676.
- Mardia, K.V., Jupp, P.E., 2000. *Directional Statistics* (second edition). Wiley.
- Tukey, J., 1980. Methodological comments focused on opportunities. In: *Multivariate Techniques in Human Communication Research* (Monge P.R., Cappella, J.N. Eds.), Academic Press, pp. 490–528.
- Watson, G.S., 1983. *Statistics on Spheres*. University of Arkansas Lecture Notes in the Mathematical Sciences, vol. 6, Wiley-Interscience.
- Watson, G.S., 1985. Interpolation of directed and undirected line data. In: *Multivariate Analysis VI* (Khrishnaiah, P.R. Ed.), Academic Press, pp. 613–625.